THE ARPANET: A USER PERSPECTIVE

by

Keith Uncapher - Information Sciences Institute

and

Vinton Cerf - Stanford University

HISTORY

In the mid-1960's, it became apparent that conventional computer communication systems were not adequate to serve the needs of the U.S. military. What was needed was a network which could survive partial destruction without total disruption of communication. Reliability several orders of magnitude better than conventional switched communication systems was needed, as well as data-security throughout the network. The network had to be modularly expandable with ease and reasonable cost.

In a 1964 series of reports entitled "On Distributed Communication" /1/, Paul Baran recommended a totally digital network based on a store-and-forward packet switching system. The rationale for this approach followed from projected decreasing costs of integrated circuits and increasing need for national data communication facilities.

Minicomputers could be distributed geographically to act as switching nodes, and would be redundantly connected by available transmission links. Data entering this minicomputer network would be formatted into packets carrying address headers and other control information. Dynamic alternate routing would be used to deliver packets to their destinations, and retransmission used to recover from transient errors detected by the receiver through checksum computations. Even speech could be transmitted through the network by first digitizing it. Simple STAR shaped networks were rejected in favor of multiply connected networks having better survivability characteristics.

The key design criteria for such a network were:

1) any transmission media could be used (copper wire, microwave, satellite, etc.)

2) each node would be connected to at least 2.2 other nodes on the average.

3) encryption-decryption would be automatic and the method used would be publicly known?

4) reliability of transmission would be improved by several orders of magnitude over conventional methods.

5) the use of very low cost electronic components would be stressed.

In 1966, an experiment was performed between the TX-2 computer at Lincoln Laboratories in Lexington, Massachusetts and the AFNSQ-7 computer at System Development Corporation in Santa Monica. The motivation was to determine whether the users of

these two systems could easily share their resources. In principle, the intent was to expand the notion of CPU time-sharing to other resource sharing as well. /2/ The experiment was only partly successful, and it became apparent that the cost of conventional leased lines, especially the very high bandwidth ones needed to meet the response time requirements of remote timesharing, was prohibitive.

In 1966, D. Davies described plans for a computer communication network similar to the design proposed by Baran earlier. Both designs had much in common with message switched systems of earlier years, but response time was far more critical in the new systems than in the old /e.g. SITA, AUTODIN; references 3 and 4/. Davies' proposal was eventually implemented at the National Physical Laboratory in England /references 5 and 6/.

By 1969, Dr. Larry Roberts, Then Chief Scientist of the U.S. Advanced Research Projects Agency, convinced ARPA and a very small segment of the U.S. computer community to embark on the design and implementation of a U.S. experimental packet swithed network, now called the ARPANET. /reference 7/. Growing from an initial 4 node network to its current nearly 50 nodes, the network rapidly involved researchers at ARPA sponsored computer laboratories and other government centers as well.

HARDWARE COMPONENTS OF THE ARPANET

The ARPANET can be divided into two logically distinct parts:

The packet switched communications subnet

The set of Hosts (serving computers) and terminal handling computers which offer or give access to network services, respectively.

The packet switched network is made up of IMPs /Interface Message Processors/ which are modified Honeywell DDP-316 or 516 computers. These are suplied by Bolt, Beranek and Newman, Inc. which maintains and operates the ARPANET for ARPA. The IMPs are connected together by 50KB telephone lines leased from AT&T by the U.S. military. Some IMPs are connected by 230.4KB lines and some by satellite links /e.g. NORSAR to SDAC and AMES to HAWAII; see Figure 1/. The satellite link from NORSAR to the U.S. is run at 7.2KB and the ground link from NORSAR to LONDON is run at 9.6KB/.

The IMPs serve as the packet switching nodes in the network. A related piece of equipment is the TIP /Terminal Interface Processor/ which is an IMP with a terminal multi-line controller and special software permitting it to connect both host computers and user terminals into the network. Logically

speaking, the TIP is really an IMP with a "fake" host inside, so TIPs are treated as if they are simultaneously hosts and IMPs./reference 8/.

Two newer developments are the PLURIBUS multiprocessor system designed by BBN to replace the current IMPs and TIPs, and the Satellite IMP /SIMP/. PLURIBUS is made up of many Lockheed SUe computers programmed to behave like IMPs or TIPs /reference 9/. The SIMP uses a multiaccess satellite channel to communicate with neighboring SIMPs /references 10-12/. These two new devices have not yet been fully integrated into the network, but are operational in a test cell at BBN. The multiprocessor system is also called a high speed modular IMP or HSMIMP, and has been attached to a spur on the network for test purposes.

Host computers are connected to the network via special purpose interfaces to the IMPs. Figure 1 shows the range of host machines connected to the network and also illustrates the heavy predominance of Digital Equipment Corporation machines. This bias is largely the result of the heavy emphasis on timesharing at ARPA research sites. DEC PDP-10 systems have a long history of providing reliable timesharing service. On the ARPANET, many of these machines are actually PDP-10X machines, having been modified with a virtual paging box built by BBN. These systems run the TENEX operating system, designed and implemented also by BBN. There are also very large scale computing resources accessible on the net including the ILLIAC IV, an IBM 360/91 at the Univercity of California at Los Angeles and an IBM 360/195 at Rutherford High Energy Laboratory at the University College London. Smaller IBM machines on the net include several 360/67's, a 360/44, and some 370's. A Burroughs B6700 is connected at the University of California at San Diego, and there are plans to connect a Control Data 7600 at Lawrence Berkeley Laboratory in California.

A number of DEC PDP-11 computers are connected to the network for the purpose of offering more elaborate terminal and special equipment access to the network than can be provided by the TIPs. In particular, experiments are underway with compressed, encrypted speech transmission using these PDP-11's for network access.

SOFTWARE COMPONENTS OF THE ARPANET

As seen from the packet switching network, the ARPANET software consists of a hierarchy of layers, like an onion. The innermost layer of software is in the packet switching network itself and makes up the protocol by which IMPs intercommunicate. At the next level is the protocol by which IMPs and hosts exchange information, and this layer is documented in (reference 13). The next level is within the hosts themselves and consists of software called the Network Control Program (NCP). Other layers can exist, either atop the NCP layer, or adjacent to it, sharing the support of the HOST/IMP protocol.

As examples of software which uses the NCP level as a basis, we can cite the File Transfer Protocol (FTP) and the TELNET protocol (which permits terminals to gain access to host services). An experimental protocol for voice transmission (Network Voice Protocol) is based on the HOST/IMP protocol and not the NCP. A protocol for host to host communication between packet switching networks (Internet Protocol) is similarly based on the HOST/IMP protocol. A Network Graphics Protocol has been proposed, but is not yet widely implemented.

A good reference to the philosophy behind ARPANET protocols is (14).

Generally, the layers of software are independent of each other, except at the interface between them. Thus, the protocol between host and IMP has remained relatively unchanged since its original specification, while the underlying IMP/IMP protocol has been very much revised. Of course, this layering may introduce some extra delay as the price paid for flexibility.

TERMINAL ACCESS TO THE ARPANET

To a user at a terminal, the ARPANET looks like a collection of distinct, interconnected computing systems. There are many ways in which a terminal can access ARPANET resources, but all of them follow the same general model:

Terminals are connected to hosts by hardwire or modem and software in the host mediates terminal access to the rest of the net.

The TIP is a special case in which the host is just software running in an IMP. Other special purpose terminal handling hosts include the PDP-11 ARPANET Terminal System (ANTS) developed at the University of Illinois in the Center for Advanced Computation and the PDP-11 ELF system developed at the Speech Communication Research Laboratory, in Santa

Barbara, California. Most other hosts on the network have software which permits terminals to connect to other hosts on the network.

To simplify server host programming for terminal access, a Network Virtual Terminal has been defined . An NVT is essentially an upper and lower case ASCII 8 bit terminal. The controls and character set of any particular terminal are translated into their NVT analogs before transmission through the network according to the TELNET protocol. The more complex programmable CRT terminaps have not yet been included in the NVT concept, although the use of such a terminal as an ASCII character mode terminal is naturally covered by the existing specification for NVT.

The user command languages offered by the TIP, ANTS, ELF, and other hosts on the network are generally different and for the peripatetic user, this can cause some confusion. Thus far in the development of the ARPANET, the preservation of the heterogeneous nature of the serving resources has been common, although there are some important exceptions which we will mention later.

FILE TRANSFER FACILITIES

Users often need to move files from one host to another. Usually these are text files, but sometimes executable program files or other binary files (e.g. data) are transferred. To meet this need, a Fle Transfer Protocol (FTP) has been designed which, under user control, permits files to be sent and retrieved between hosts.

One part of the FTP facility permits "network mail" to be shipped to well known mailboxes at serving hosts. Special programs which permit users to compose messages or examine them make use of the FTP to transmit messages among users. In a distributed environment, geographically separated groups can cooperate with each other on common software or hardware projects, keeping up to date via the message facility. It is very likely that the greatest impact that the network has had on the ARPANET research community is this vastly improved communication mechanism. Although there are no published statistics on the fraction of time the network spends moving messages around, one of the authors spends a large proportion of his network connect time either composing or examining electronic mail. The facility is easily integrated into a normal working environment, and in fact, it becomes addictive.

REMOTE JOB ENTRY FACILITY

Some of the service sites on the ARPANET, especially the IBM systems, offer batch facilities. A Remote Job Entry protocol has been defined which uses the File Transfer Protocol to move

input and output files from user to server hosts. Through the RJE facility, a user can submit a file to be entered into the serving host's job stream and indicate where the output file(s) should be deposited after the job has completed.

RESOURCES AND THEIR USE

Because the ARPANET is strongly oriented toward interactive use, it is no surprise that there are a large number of different text editing systems on the network. Coupled with the FTP mechanism, this permits users to compose text (papers, documentation, program source, etc.) using their preferred editors and then ship the resulting files elsewhere, if need be, for further processing. As an example, this paper was written using the NLS editing system at Stanford Research Institute's Augmentation Research Center, but it was printed at Stanford University on an upper/lower case printer.

An appendix has been provided with example uses of the network in all their detail, for the curious minded. In the following few paragraphs, we list some typical ARPANET uses, but leave out the details of actual interaction between user and serving system.

Electronic Mail Facility

In all of the TENEX systems, and in most of the other hosts attached to the network, a program which will accept a message from a user and transmit it to a specified receiving host where the message will be held until the recipient asks to see it. Often the serving host systems will alert a user that he has "mail" waiting when he first logs into the server.

Users typically make use of this facility to exchange short technical notes or to leave messages which would otherwise have had to go by regular mail, with the associated delay. Creating such an environment for fast communication makes it possible for research teams in widely distributed locations to cooperate intimately on a common research goal. The Speech Understanding Research project, sponsored by ARPA, has been able to achieve fusion of effort, largely through the use of this facility. An experiment in the design of protocols for connecting hosts in different works involves three main groups in London, Stanford, California, and Boston. The mail facility has been indispensable.

An extension to the mail facility permits text files to be inserted into outgoing mail messages, thus permitting more formal and longer documents to be exchanged. Another typical technique is to send a short message describing the location of a file to be examined by a group of researchers. The participants can then independently read

and comment on this material. A number of conference and journal papers have been cooperatively written by authors thousands of miles apart, using this facility.

TENEX "Link" Facility

In TENEX, and in several other systems on the network, it is common to find a mechanism for linking two or more terminals together so that whatever is typed on one is typed on all. This facility is of tremendous value as a teaching aid, since it is possible to give a demonstration of the use of a particular program or system to a collection of observers whoo see precisely how to use it. This aid is particularly valuable if a question comes up on the use of a system which is maintained and running at a site very distant from the using site. If there are not too many time zones involved, it is often possible to track down the responsible programmer and get him or her to answer the question on line. Of course, this facility can also be used as a crude sort of "forum" but for more than 2 participants, it can become rather confusing.

FORUM - A Conferencing Facility

The FORUM program was conceived and written by researchers at the Institute for the Future in Menlo Park. Essentially, it offers a way of recording and organizing the ideas of many people on a subject or set of subjects. Comments from the participants are kept in chronological order, but can be withdrawn and listed from a data file based on a variety of conditional sesearch commands (e.g. all comments during a particular period by a given set of authors, etc.). Although this facility is not a panacea, it serves well in the absence of full bandwidth video/audio, and also provides a machine readable record of the discussion which is highly useful. One can potentially discover where, in the course of an investigation, it went astray, or what series of considerations led to a decision. Usually this information is lost in everyone's hazy memory of history.

Cross-compilation and Assembly

An increasingly common use of network computing capability shows up in the support and maintenance of minicomputers attached to the ARPANET. It is often simpler and more efficient to build compilers or assemblers which will run, for example, on TENEX, but produce programs which can only be run on a particular minicomputer. The File Transfer Protocol can be used to move the executable files to the desired destination. Nearly all of the software for the two "front-end" systems, ELF and ANTS, have been produced and are maintained on TENEX and the Burroughs B6700 respectively. Documentation is often produced on yet

another system, for example the NLS facility at Stanford Research Institute.

Multi-host Systems

Another example of specialized use of different network resources shows up in the REDUCE system which normally runs in interactive mode on a TENEX at the University of UTAH. This is a symbolic arithmetic facility which permits a user to manipulate complex equations with the help of a computer. Symbolic differentiation, matrix inversion, and so on are available. Once a set of equations has been satisfactorily generated, it is possible for REDUCE to automatically generate a FORTRAN program which can numerically solve the set of equations. This program can be sent to the 360/91 at UCLA, where it is executed, and the results returned via the RJE protocol to the user at UTAH (or where ever he happens to be).

A similar facility is available at MIT's Mathematics Laboratory where the MACSYMA system is in operation. Some experiments have been performed using MACSYMA as a "subroutine" to a program running at MIT'S Dynamic Modelling System. These are just the first steps towards more intimate cooperation between processing systems on the network.

Sports and Recreation

It should be no surprise that the ARPANET probably houses the world's largest collection of computer games. Ranging from Chess to word games like JOTTO and of course, SPACEWAR, these games are the byproducts of the natural interest of the programming community which has grown up with the network. Some of the interactive facilities have potential as computer assisted instruction tools. Some offer challenging problems in communications protocol design. It often happens that these games must be banned during peak hours since they can absorb a great deal of computing power when it is needed elsewhere.

On-Line News Service

At Stanford University's Artificial Intelligence Laboratory, low speed lines from the Associated Press news system and the New York Times news service are connected to a PDP-10/PDP-6 dual processor system. A sophisticated program accepts the news stories, concatenates the various "takes" into a series of story parts, and also forms an inverted file of keywords so that the database (about the last two weeks of news) can be selectively searched for stories of special interest. It is not hard to imagine that

this facility will someday represent the normal delivery mechanism in the home.

Caveat Emptor

One of the more troublesome aspects of the network is in fact the tremendous number of services available in the nooks and crannies of computers everywhere. A user often has no idea where to look for a given service, or once found, is not always sure how to use it. The first problem, finding the resources, has been partly solved through the publication of a "Resource Notebook." The big problem here is getting the authors of programming facilities to document them sufficiently to make it possible for an innocent user to learn how to access the resource. The second problem is solved more and more often through the use of on-line help. When a user is unsure what to do next, he or she can typically type "?" or "HELP" and get a useful response. This kind of "hand-holding" will be essential if networking is ever to attract the relatively unsophisticated user.

A SERVER VIEWPOINT

One of the primary goals of the ARPANET is resource sharing. In simplest form, this means the availability of a host providing a set of services to users throughout the network.

Owing to the simple access to all the network resources, it is clear that a service site is rather like a "fishbowl." That is, the character of the service, the reliability and the general quality of it are very apparent to any user on the network. This has had a strong upgrading effect on the quality and performance of services offered on the net; it is difficult to sell poor performance, and the user community knows which sites are good and which are not. For the university environment, it is a new challenge and one not necessarily met with proper response. In general, the fish bowl concept has improved the reliability, availability, and utility of service sites very substantially.

Very much like global commercial time sharing, a service site like ISI in Southern California has a prime shift which is proportional in length to the number of time zones in which frequent users reside. For ISI, it is 0500 to 2100 (Boston to Honolulu) with occasional usage from London.

As resources on the network become more useful, the users' expectations seem to demand telephone-like reliability. Hence, 24 hour per day, 7 days per week continuous service seems essential, and it is a reflection of the growing importance and dependency on such services by many users (i.e. it becomes addictive...).

All of this says a great deal about the importance of the host site reliability and management philosophy, both of which must be totally supportive of excellent service. One should note in passing that with the variety of services available on the network, there is a tendency for more and more users to spend hours per day accessing services. Hence, the dependency on the network's rsources is vital indeed and availability of reliable services become paramount as users shape their work and lives around on-line interactive use. Along these lines, response time from interactive systems becomes a very serious issue. Once a dependency is established on sophisticated text editing services, for instance, the time gained from these services will be lost utterly if response time slows the user down below ordinary typewriter speeds. Nobody likes a sluggish typewriter, and an overloaded service site can easily consume twice as much of a user's time as it should.

Host sites on the network range from commercial time-sharing to very esoteric data base facilities and languages only usable by the skilled computer scientist (or the well-versed undergraduate hacker). There is a tendency toward more uniform service and the notion of a standard command language which makes all the heterogeneous resources on the net appear to be uniformly accessible is popular indeed. Some steps in this direction have been taken by the BBN TENEX design team in the form of a Resource Sharing Executive experimental monitor (RSEXEC, reference 15). The popularity of the TENEX systems is in part a reflection of the great breadth of languages and services provided and the desire for uniform access methods. It also reflects the fact that there is little financial or emotional support for private monitor development in the U.S.; there is simply no way to compete with the tremendous resources of the manufacturers in this domain. It is worth mentioning, however, that both TENEX and MULTICS operating systems were developed in a university/industry cooperative effort.

Another trend among host sites is functional specialization. With low cost, on-demand service independent of physical location via the ARPANET, the following represent typical examples:

a) a few hosts specialize in number crunching, i.e. the Burroughs 6700, IBM 360/91, 360/195, and ILLIAC IV.

b) many hosts offer standard TENEX.

c) the UNICOM trillion bit store at Ames Laboratory

d) a special purpose front end processor for image processing research. One machine in California will service the needs of three image research centers.

e) the DATACOMPUTER at Computer Corporation of America, specializing in data base applications.

SUMMARY

The ARPANET experiment has uncovered some important and dramatic results, among which are:

a) researchers can function well using remote services and interacting closely while still geographically distributed.

b) service sites on the network become highly visible with respect to performance and must "shape up" to maintain the their clientele.

c) the widely disbursed and large user community makes it possible to develop special services with reasonable expectation of sufficient demand.

d) users can and do access special services that are remote simply because they are easily accessible, and the cost is low enough to make it worthwhile.

e) server sites must supply stable service on a 24 hour per day, 7 days per week basis, with virtually assured access and reasonable response to interactive requests.

f) great economic gain and, in general, improved service will result when two or more machines are installed at a service site; there is an economy of scale.

g) when worldwide netting of resources is available, time zone differences can make dramatic cost reductions by offering "prime shift" service at off-hour times on a machine for various sets of users around the world. Some evidence of this is already present in the TYMSHARE Corporation TYMNET and in the General Electric MARK III Network.

h) a small number of research centers now function well without local processing capability and hence are totally dependent on ARPANET remote resources. Perhaps it is fair to say that with adequate intelligence in a terminal or a terminal support processor, the network makes it possible to work as if all the service sites are virtually in your own backyard.

This short report has only scratched the surface of a vast area of technical, social, and economic interest. Two exceptionally rich sources of primary material and bibliographic references are found in /references 16 and 17/ and we recommend them to the interested reader.

Postel, "Function Oriented Protocols for the ARPA Computer Network," Proceedings SJCC 1972, p. 271-279.

15. Thomas, R., "A Resource Sharing Executive for the ARPANET, " Proceedings NCC 1973, p. 155-163.

16. Chu, W.W. (editor) Advances in Computer Communications, ARTECH House Inc., 1974 (reprints of many important articles in computer communications).

17. McQuillan, J.M., Adaptive Routing Algorithms for Distributed Computer Networks, BBN Report No. 2831, May 1974 (also the author's PH.D. Dissertation at Harvard University).

APPENDIX - EXAMPLES OF USAGE

DIALING TO A TIP

(When the TIP answers the call, the user types "E" to identify the terminal type and speed, and the TIP responds with:)

HELLO 322#: 22

(This means TIP program version number 322 is running and you are on line number 22)

(You then select a serving host by typing its "call number")

@L 86

(In this case, the "@" is an escape character which tells the TIP to interpret what follows as a command, rather than sending it on to a serving computer in the network. Host 86 is the PDP-10 at Information Sciences Institute.)

Logger (The TIP tells you it is trying to connect to ISI)

Open (The TIP tells you the connection is open to ISI)

Message slots are now being allocated.

Type LOG or GLOG; type OFFQUOTA for more information.

(This is from the ISI machine; it asks that you login -- i.e., identify yourself.)

ISI-KA-TENEX 1.32.8, ISI-TENEX EXEC 1.51.4

(The ISI PDP-10 identifies itself and which versions of

its operating system and command language executive are running./

@LOGIN CERF

/the "@" character is a prompt from the ISI PDP-10 to tell you it is listening. You identify yourself./

(PASSWORD)

/TENEX is asking for your secret password. It will not print out on your terminal when you type it./

(ACCOUNT #) 1

/You also type the account number to which this session should be charged./

JOB 25 on TTY6 17-OCT-74 09:22

/Tenex assigns you to teletype number 6 and job number 25 and tells you what day and time it is./

TENEX WILL GO DOWN THU 10-17-74 2345 TIL FRI 10-18-74 0500

/TENEX advises you that it is scheduled to be out of service from 11:45 PM until 5 AM the next day./

YOU HAVE A MESSAGE

TENEX warns you that there is "electronic mail" waiting for you in your mailbox./

READING ELECTRONIC MAIL

@RD

/You decide to read your mail./

TECO 1.29

/"RD" is a program based on version 29 of the TECO editor./

312 CHARS

/RD says there are 312 characters of text waiting in your mailbox./

1 17 OCT 1974 CERF TEST OF JOHANNESBURG LINK

/Message number 1 is dated 17 October, came from CERF and is about the "Test of Johannesburg Link"/

TYPE MH$ FOR HELP

(RD lets you know how to get help.)

*

(RD is now prompting you for your command.)

*1MT$

(This command says type message number 1. The "$" is the way TENEX echoes the ESCAPE character, since ESCAPE is not a printing character.)

Net mail from site USC-ISI rcvd at 17-OCT-74 10:05:25

Date: 17 OCT 1974 1005-PDT

From: CERF at USC-ISI

(Message originated from user CERF at the ISI system.)

Subject: TEST OF JOHANNESBURG LINK

To: CERF at USC-ISI

cc: VANDERVEER

(A copy was sent to user VANDERVEER implicitly at USC-ISI.)

GERT:

I HOPE THAT DON MALPASS WAS ABLE TO HELP YOU GET THE MODEM AND TERMINAL WORKING. IF HE HAS, PLEASE SEND A MESSAGE TO ME AT ISI. THANKS, AND GOOD LUCK.

VINT
